



ELSEVIER

Physica A 299 (2001) 234–246

PHYSICA A

www.elsevier.com/locate/physa

Price fluctuations from the order book perspective—empirical facts and a simple model

Sergei Maslov^{a,*}, Mark Mills^b

^a*Department of Physics, Brookhaven National Laboratory, Upton, NY 11973, USA*

^b*1320 Prudential, Suite 102, Dallas, TX 75231, USA*

Abstract

Statistical properties of an order book and the effect they have on price dynamics were studied using the high-frequency NASDAQ Level II data. It was observed that the size distribution of marketable orders (transaction sizes) has power law tails with an exponent $1 + \mu_{\text{market}} = 2.4 \pm 0.1$. The distribution of limit order (or quote) sizes was found to be consistent with a power law with an exponent close to 2. A somewhat better fit to this distribution was obtained by using a log-normal distribution which has an effective power law exponent equal to 2 in the middle of the observed range. The depth of the order book measured as a price impact of a hypothetical large market order was observed to be a non-linear function of its size. A large imbalance in the number of limit orders placed at bid and ask sides of the book was shown to lead to a predictable short term price change, which is in accord with the law of supply and demand. © 2001 Elsevier Science B.V. All rights reserved.

PACS: 89.65.Gh; 89.75.Da; 89.75.Fb; 05.40.Ca

Keywords: Limit order; Order book; Price fluctuations; High-frequency data

As a result of collective efforts by many authors the list of basic “stylized” empirical facts about market price fluctuations has now begun to emerge.¹ It became known that the histogram of short term changes in price $\delta p(t) = p(t + \delta t) - p(t)$ has “fat” power-law tails: $\text{Prob}(\delta p > x) \sim x^{-\alpha}$. The exponent α was measured to be close to 3 in major US stock markets [1] as well as foreign exchange markets [2–4, see Tables 4 and 5]. The other well established empirical fact is that while the sign of $\delta p(t)$ measured at different times has only short term correlations, its magnitude $|\delta p(t)|$ (or alternatively its square $\delta p(t)^2$) has a long term memory as manifested by slowly

* Corresponding author. Tel.: +1-631-344-3742; fax: +1-631-344-2918.

E-mail addresses: maslov@bnl.gov (S. Maslov), mmills@htcomp.net (M. Mills).

¹ For an up to day references to econophysics literature see articles in these proceedings.

decaying correlations. The correlation function was successfully fitted by a power law $t^{-\gamma}$ with a small exponent $\gamma \simeq 0.3$ [5,6] over a rather broad range of times.

Several simplified market models were introduced in an attempt to reproduce and explain this set of empirical facts [7–12]. The current consensus among econophysicists seems to be that these facts are a manifestation of some kind of strategy herding effect, in which many traders lock into the same pattern of behavior. Large price fluctuations are then explained as a market impact of this coherent collective trading behavior. Any model aiming at understanding price fluctuations needs to define a mechanism for the formation of the price. Here the usual approach is to postulate some empirical (linear or non-linear) market impact function, which reduces calculating prices to knowing the imbalance between the supply of and the demand for the stock at any given time step.

Recently, one of us (S.M.) has introduced a toy model [13] in which the same standard set of stylized facts, albeit with somewhat different critical exponents, was generated in the absence of any strategic behavior on the part of traders. The model uses a rather realistic order-book-based mechanism of price formation, which does not rely on any postulated market impact function. Instead, price fluctuations arise naturally as a result of changes in the balance of orders in the order book. The long memory of individual entries in this book gives rise to fat-tailed price distributions and volatility clustering. Every market has two basic types of orders, which we would refer to as limit and market orders. A limit order to sell (buy) is an instruction to sell (buy) a specified number of shares of a given stock if its price rises above (falls below) a predefined level, which is known as the execution price of a limit order. A market order on the other hand is an instruction to immediately sell (buy) a specified number of shares at whatever price currently available at the market. Here, we do not make a distinction between a true market order and a marketable limit order, which can be immediately filled by previously placed limit orders on the complimentary side of the market, and refer to both of them as ‘market orders’. The model of Ref. [13] assumes the simplest possible mechanism for the dynamics of individual orders in the order book. At each step, a new order is submitted to the market. With equal probabilities this can be a limit order to sell, a market order to sell, a limit order to buy, or a market order to buy. All orders are of the same unit size, and a new limit order to sell (buy) is placed with a random offset Δ above (below) the most recent transaction price. In spite of its utmost simplicity the model has a surprisingly rich behavior, which up to now was understood only numerically. The distribution of price fluctuations has power law tails characterized by an exponent $\alpha=2$, while the correlation function of absolute values of price increments decays as $t^{-0.5}$.

Of course, the dynamics of a real order book is much more complicated than the rules of the toy model from Ref. [13]. First of all, in real markets, both market and limit orders come in vastly different sizes and exist for various time frames. Secondly, participants of real markets do use strategies after all. In particular, both under-capitalized speculators and well-capitalized market makers avoid static public display of their willingness to accept a given price. In practice that means that a quote or a limit order may represent only a part of a large pending order. An active market participant may

adjust a previously placed quote in response to sudden price movements, or display and later withdraw a “fake” quote, which is sufficiently far from the highest bid and lowest ask prices, so that it is almost never filled. This last strategy creates an illusion of an imbalanced order book, which can confuse other traders. Finally, there is a practically all-important matter of time delay between the actual state of the order book and whatever a particular trader observes on his/her screen. Prior to electronic data transmission, investors might not know at what price the queue is matching their buy and sell orders until long after the transaction took place. On the other hand, market makers have always had near immediate access to completed transaction data. With modern computerized markets, there is a much shorter delay between a transaction’s completion time and trader’s awareness of the event, but yet the delay still exists. The authors are familiar with day-traders moving from one city to another based on an empirical discovery of as little as a 1 s improvement in data delivery. During peak activity periods (near the open, close and significant news events), data delivery delays of 15–30 s are sometimes experienced. The inhomogeneity of those delay times for different market participants contributes to the wide variety of strategies employed by traders. It is common knowledge among professional traders that some novices will attempt to day-trade using web-based data and order entry systems. During peak activity, their data may be 5–10 min old. Professional traders anticipate this by taking a position seconds after a significant short term price change, and then exchanging this position with a novice a few minutes later.

In this work, we attempt to establish some empirical facts about the statistical properties and dynamics of publicly displayed limit orders using data collected in a real market. The purpose of this analysis is twofold. First of all, these new observations would extend a rather narrow list of stylized facts about real markets. In econophysics as in other branches of physics (or any other empirical science for that matter) the only way to choose among many competing theoretical models is to make new empirical observations. Since the high frequency data about the state of an order book is much harder to collect than the highly institutionalized record of actual transactions, to our knowledge this investigation was never attempted by members of the econophysics community before. Second, we hope that the study of a real order book’s dynamics would suggest new realistic ingredients that can be added to a toy model of Ref. [13] to improve its agreement with the extended set of stylized facts.

Markets differ from each other in precise rules of submission of orders and the transparency of the order book. In the so-called order-driven markets there are no designated market makers who are required to post orders (quotes) on both bid and ask sides of the order book. Instead, the liquidity is provided only by limit orders submitted by individual investors. Versions of this market mechanism are employed in such markets as Toronto Stock Exchange (CATS), Paris Bourse (CAC), Tokyo Stock Exchange, Helsinki Stock Exchange (HETI), Stockholm Stock Exchange (SAX), Australian Stock Exchange (ASX), Stock Exchange of Hong Kong (AMS), New Delhi and Bombay Stock Exchanges, etc. Major US markets use somewhat different systems. In the New York Stock Exchange individual orders are matched by a specialist who

does not disclose detailed data regarding the contents of his order book. This reduces the transparency (or openness) of the order book to market participants. The NASDAQ Level II screen is the closest US equivalent to an order book in an order-driven market. Since the contents and dynamics of individual entries on this screen are main subjects of the present work they will be described in greater details later on in the manuscript.

Before we proceed, we would like to make an important disclaimer regarding the terminology used in this paper. To avoid overwhelming our readers by a variety of different financial terms describing similar concepts, in this work, we would refer to any yet unfilled order present in an order book as a ‘limit order’. While this is strictly true for an order driven market, using this term to describe a market maker’s quote on the NASDAQ Level II screen may seem a bit confusing at first. However, it makes sense in this context. Indeed, both individual limit orders in an order-driven market and market maker’s quotes on the NASDAQ Level II screen can be viewed just as commitments to buy (sell) a certain number of shares at a given price should the queuing mechanism match this order with a complementary marketable order. The only detail which distinguishes a market maker at NASDAQ from a regular trader in an order-driven market is that by NASDAQ rules, the market maker must maintain both buy and sell limit orders, changing their price level and volume within domains established by exacting timing rules. But in zero order approximation one can simply forget that these two quotes come from the same source and think about them just as two individual ‘limit orders.’

The other simplification adopted in this work is that we do not make a distinction between a true market order and a marketable limit order, placed at or better than the inside bid or ask price, and refer to both of them as ‘market orders’. From this point of view, a transaction always happens when a ‘market order’ (or a marketable limit order) is matched with a previously submitted ‘limit order’ (or a quote by the market maker). The size of an individual transaction is therefore a good measure of a market order’s size in our definition.

The real time dynamics of an order book is a fascinating spectacle to watch (see e.g. www.3dstockcharts.com). For frequently traded stocks, it is in a state of a constant change. The density of limit orders goes up when more traders select to submit limit orders rather than market (or marketable) orders. In the opposite case of a temporary preponderance of market orders, the book gets noticeably thinner. In addition to these fluctuations in the density and number of limit orders, any serious imbalance in the number of limit orders to buy and limit orders to sell which are not too far from the current price level gives rise to predictability of short-term price changes. This change reflects intuitive notions regarding supply and demand. i.e., the price statistically tends to go up in response to an excess number of limit orders to buy and down in the opposite case. It is by observing all of this in real time that one understands that the balance of individual orders in the order book is the *ultimate* source of price fluctuations.

In this work, we study statistical properties of the data that one of us (M.M.) has collected, while trading on the NASDAQ market. Even though NASDAQ is a

quote-driven (dealership) market, due to reasons explained above we believe that our study should also apply to order books in order-driven markets. Indeed, many of our conclusions are remarkably similar to those reported for order-driven markets in the recent economic literature (see e.g. Refs. [14,15]). The NASDAQ Level II data for a given stock lists current bid and ask prices and volumes quoted by all market makers and Electronic Communication Networks trading this stock. For example, the line: JDSU GSCO K NAS 112.625 500 114.0625 500 can be interpreted as a display of Goldman Sachs' (GSCO) intent to buy 500 shares of JDS Uniphase Corporation (JDSU) at 112.625 per share and sell 500 shares at 114.0625 per share. Each such market maker entry usually conceals a large secondary order book of limit orders submitted to this market maker by his clients. Those “outside” bids and asks, i.e., private limit orders at price levels more distant from the publicly displayed best (or “inside”) bid or ask, generally remain hidden to most market participants. The concept of second hierarchical level in the order book at NASDAQ can perhaps be best illustrated on an example of Electronic Communication Networks (ECN) such as Island (the ECN symbol ISLD). In this case, the “hidden” book can be actually viewed (e.g. at the Island's website (www.island.com)), while the only part of this book which is visible at the NASDAQ Level II screen is its highest bid and lowest ask prices and volumes. There they are shown as any other market maker entry: JDSU ISLD O NAS 113.75 200 114 800.

In the course of one trading day we recorded ‘snapshots’ of the order book for one particular stock at time intervals which are on average 3 s apart. We were unable to exactly account for the network delay between the time when a particular quote was issued by the NASDAQ order-matching queue and the time when it was received by us (i.e., our time stamp). Most of the time this delay is less than a second, but it is known to occasionally exceed 15 or even 30 s during the times of especially heavy trading volume. This record was subsequently binned by the price. Prices and aggregate volumes at four highest bids and lowest asks were kept in the file. Due to the discreteness of stock price at NASDAQ several market makers are likely to put their quotes at exactly the same price. In our file we kept only the aggregate volume at a given price, equal to the sum of individual limit orders (quotes) by several market makers. A file collected during a typical trading day contains on average 7000 time points.

The first question we addressed using this data set was: what is the size distribution of limit and market orders? In Fig. 1, we show the cumulative distribution of *market* (marketable) order sizes (or alternatively the sizes of individual transactions) calculated for all stocks and trading days for which we have collected the data. From our record we know only the total number of traded shares and the total number of transactions which occurred between the two subsequent snapshots of the order book. This average number of transactions per snapshot varies between 3 and 5.5 for different stocks in our data set. The size of a market order used in Fig. 1 was defined simply as the change in the traded volume divided by the (usually small) number of transactions that occurred between two subsequent snapshots of the screen. All our data are consistent with transaction (market order) sizes being distributed according to a power law $P(x) \sim x^{-1-\mu_{\text{market}}}$ with an exponent $\mu_{\text{market}} = 1.4 \pm 0.1$. Authors of Ref. [16] have analyzed

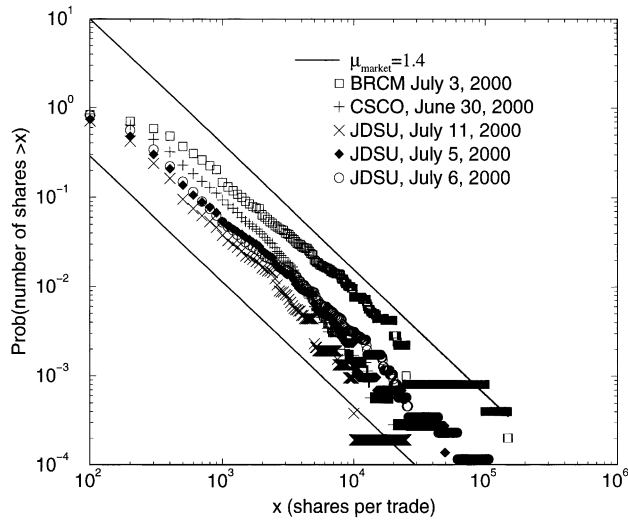


Fig. 1. The cumulative distribution of market order (transaction) sizes. The straight line has the slope $\mu_{\text{market}} = 1.4$.

the distribution of volumes of individual transactions for largest 1000 stocks traded at major US stock markets and arrived at a similar *average* value for the exponent $\mu_{\text{market}} = 1.53 \pm 0.07$ (ξ in their notation). They also plotted the histogram of this exponent measured for different individual stocks (see Fig. 3(b) in Ref. [16]), showing substantial variations.

The distribution of *limit order* sizes, to our knowledge, was never analyzed in the literature before. To make the histogram of this distribution we used sizes of limit orders at a particular level in the order book from all snapshots made throughout one trading day. We found that this histogram can also be approximately described by a power law form. The data for different levels of bid and ask prices (level 1 being the highest bid and the lowest ask) for two of our stocks are presented in Figs. 2 and 3. In both cases, all distributions were found to be consistent with an exponent $\mu_{\text{limit}} = 1.0 \pm 0.3$. The quality of the power law fit is rather poor though with error bars on μ_{limit} in each *individual data set* around ± 0.1 , while it was close to ± 0.015 for the power law fits in Fig. 1. The above error bars of ± 0.3 also include fluctuations of the exponent between different data sets. Looking for a better fit to our limit order data we repeated the above analysis using cumulative histograms and we saw that a log-normal distribution approximates our data over a wider region (see Fig. 4). The best fit to a log-normal distribution has similar parameters for different stocks, trading days, and levels in the order book. The best empirical formula for the probability distribution of limit order sizes is thus $P(x) = x^{-1} \exp(-(A - \ln(x))^2/B)$, with parameters A and B fluctuating around 7 ± 0.5 and 4 ± 0.5 in all of our data sets. This formula indeed gives the effective power law exponent $\mu_{\text{limit}} = 1$ for $x \simeq 8000$ i.e., near the center of our range.

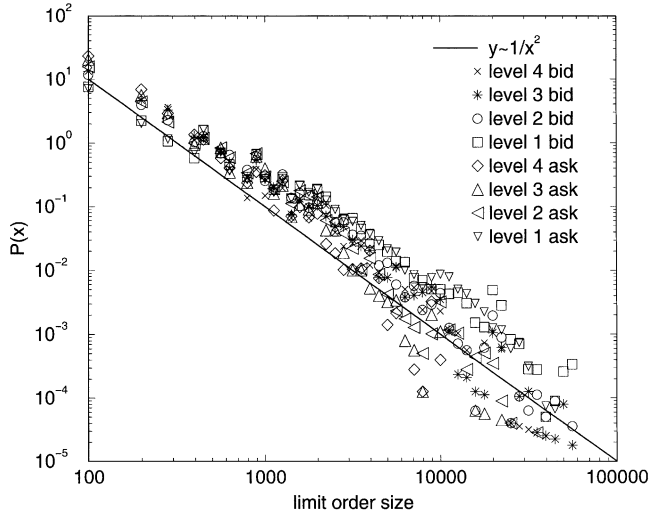


Fig. 2. The size distribution of limit orders (consolidated quotes) placed at different levels in the order book for the stock of the JDS Uniphase Corporation (ticker symbol JDSU) traded on July 5, 2000. The straight line has the slope $1 + \mu_{limit} = 2$.

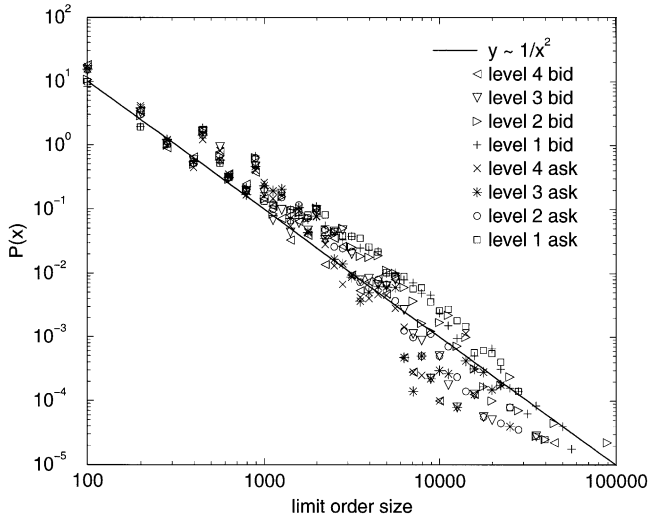


Fig. 3. The size distribution of limit orders (consolidated quotes) placed at different levels in the order book for the stock of the Broadcom Corporation (ticker symbol BRCM) traded on July 3, 2000. The straight line has the slope $1 + \mu_{limit} = 2$.

We next concentrate on calculating the depth of the order book at any given bid and ask level. The depth of the order book is an important measure of the liquidity of the market for a particular stock. For a given state of the order book one can measure the total volume (number of shares) $N(\Delta p)$ of limit orders with execution prices lying

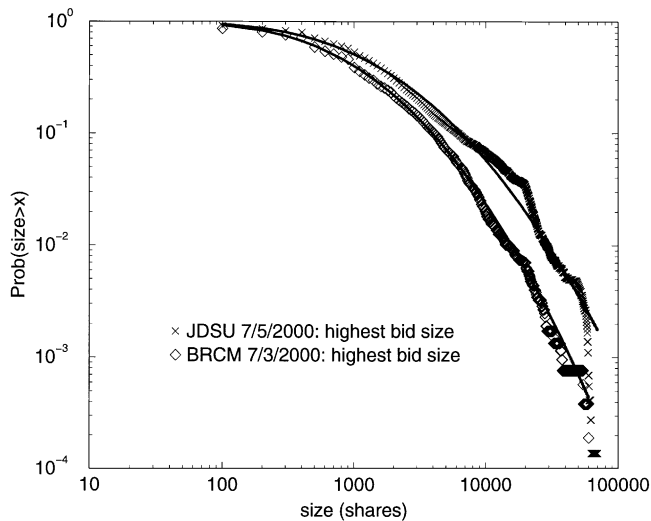


Fig. 4. The cumulative distribution of highest (level 1) bid sizes. Solid lines are best fits with the cumulative histogram of a log-normal distribution $P(x) = x^{-1} \exp(-(A - \ln(x))^2/B)$. The best fit parameters are $A_1 = 6.94$, $B_1 = 4.20$ for the JDSU, and $A_2 = 6.57$, $B_2 = 3.56$ for the BRCM.

within a certain price range Δp from the middle of the highest bid/lowest ask spread. The function $\Delta p(N)$, which is the functional inverse of $N(\Delta p)$ can be thought of as the *virtual* impact that a hypothetical market order of volume N would have on the price of the stock. That is to say, a hypothetical trader willing to immediately sell N shares would have changed the price by $\Delta p(N)$, provided that no new limit orders (quotes) would appear, while his order is executed. It is important to emphasize the word *virtual* here. Indeed, in real markets new limit orders would be immediately submitted by market makers (or speculators in order-driven markets) in response to the arrival of a large market order. The first step in quantifying the depth of the limit order book is to measure the average price difference between different levels of the book, e.g. the average gap between the prices of the highest and the next highest bids. For both bid and ask sides of the book at all levels, the average price gap between levels was measured to be around \$0.064 for the CSCO stock traded on June 30, 2000, \$0.08 for the JDSU stock traded on July 5, 2000 and \$0.12 for the BRCM stock traded on July 3, 2000. At the same time, the number of outstanding shares of the above three companies were 7.3, 1.3, and 0.24 billion shares correspondingly. Also, the average daily volume of CSCO on the particular day when our data were recorded was larger than that of JDSU and BRCM by approximately a factor 1.5. One can see that more capitalized and more frequently traded companies have on average smaller gaps between their consecutive bid and ask prices. In case of JDSU and BRCM, the average bid–ask spread (i.e., the difference between the lowest ask and the highest bid prices) was measured to be some 10–20% smaller than the average gap between two levels on the same side of the book. Also, for both these sets, gaps on the ask (limit

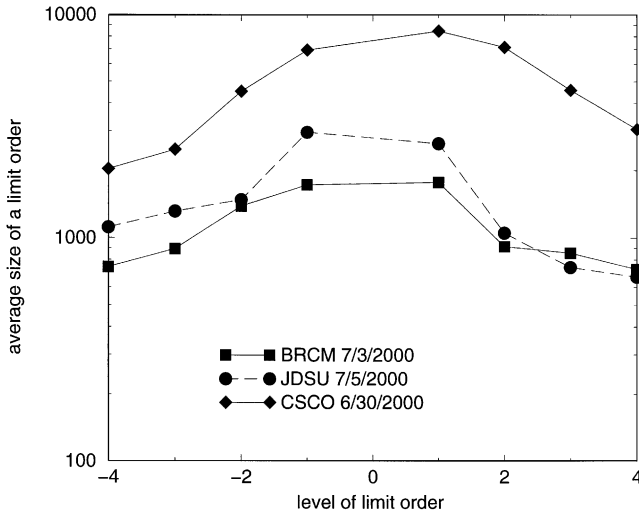


Fig. 5. The average number of shares offered in limit orders at a given level as a function of the level number. Negative levels correspond to limit orders to buy (bids), while positive to limit orders to sell (asks).

orders to sell) side seem to be some 5–10% higher than on the bid (limit orders to buy) side. It is not clear if that was just an artifact of the trading day or a sign of some real asymmetry. More interesting behavior was observed for the average size of a limit order as a function of level of the order book. The average size reaches its maximum at level 1 of the book (highest bid/lowest ask) and gradually falls off with the level number (see Fig. 5). Using the data for the average volume at each level and the average price difference between levels one easily reconstructs the curve $\Delta p(N)$ (see Fig. 6). From this curve one concludes that the *virtual* price impact Δp of a market order is a nonlinear function of the order size N . Similar results were observed for the limit order book at the Stockholm Stock Exchange by Niemeyer and Sandas (see Fig. 8 in Ref. [17]). To derive a concise formula for $\Delta p(N)$ we fit it to the power law $|\Delta p(N)| \sim |N|^\delta$ separately on positive and negative sides of each curve. The exponent δ of this fit was measured to be 2.05 ± 0.05 for four out of six curves, while in the remaining two it was as high as 2.5 and 2.7. In Ref. [18] it was argued that the price impact function should have an exponent $\delta = 0.5$. This conjecture was later used in several models to arrive at the empirically observed value of the exponent α of the fat tails in the histogram of price fluctuations. Our virtual market impact function characterized by $\delta > 2$ has the opposite convexity compared to that with $\delta = 0.5$. We attribute this discrepancy to the difference between virtual and real market impacts, where the latter is dramatically softened by actions of speculators.

The subject of speculators brings us directly to the last question we addressed using our data: can one use the information contained in the order book to predict the magnitude and direction of price changes in the near future? Many seasoned day traders would answer yes to this question. From the law of supply and demand, one expects

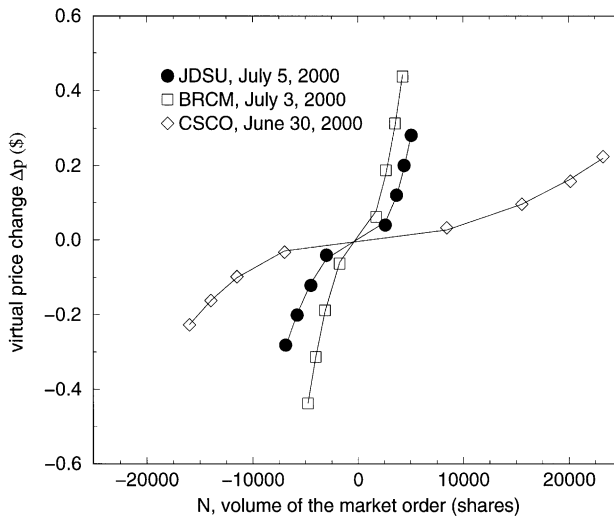


Fig. 6. The virtual impact of a market order calculated from the density of limit orders in the order book. Negative x corresponds to market orders to sell, while positive—to market orders to buy. Solid lines are power law fits to the data performed separately on positive and negative sides. The exponent δ of the best fit was measured to be 2.05 ± 0.05 except for the negative part of the CSCO curve (2.5) and the positive part of the JDSU curve (2.7).

that a significant excess of limit orders to sell above limit orders to buy (excess supply of stock) would push the price down while in the opposite case, the price would go up. It means that a speculator who has access to the current state of the order book can predict (and use this prediction for his/her profit) the direction of price change in the near future. The first way to measure the short term predictability of market price changes using our data is to concentrate on those moments in time when the total number of shares contained in limit orders to sell and limit orders to buy differ by a significant number of shares. In principle, this amount should be selected proportional to the average daily volume of transactions for each particular stock, yet in our calculations we fixed it to be 10,000 shares for each of the stocks in our data sets. Also, we looked only at the imbalance between volumes offered at highest bid and lowest ask prices. We then averaged price increments at times immediately following the moment of large excess demand (or supply) over all events when this excess occurred and plotted it as a function of time in Fig. 7. From this plot one can conclude that indeed as expected from the law of supply and demand, an excess of limit orders to buy drives the price up, while an excess of limit orders to sell drives it down. In our data set this predictability of future prices lasts for only a few minutes (only for 30 s for some of the stocks). Therefore, speculators who want to use this effect need to act quickly and to have a very fast and reliable connection to the main computers at NASDAQ. Yet another way to visualize the effect the imbalance of supply and demand has on the price is to calculate the average change in price of the stock during a fixed time interval Δt conditioned at a certain value of the

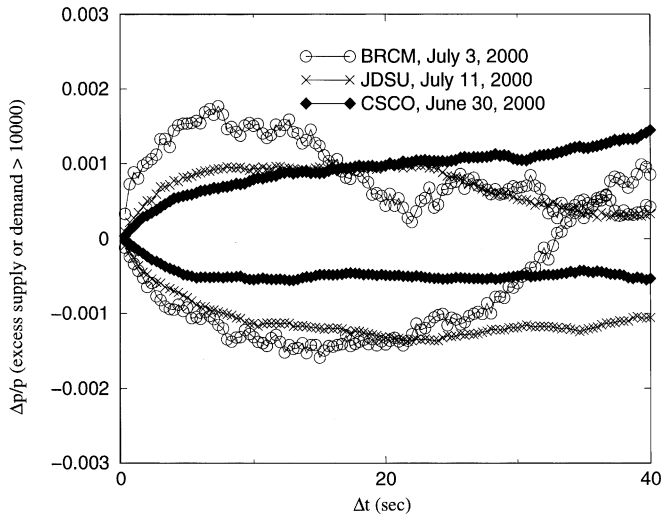


Fig. 7. The market impact of a large imbalance (larger than 10,000 shares) in the number of shares offered at the highest bid and lowest ask prices. The upper portions of curves correspond to the excess demand for the stock, while lower ones for the excess supply.

imbalance of the order book before the change. In Fig. 8 we plot the average 1-min price change² as a function of the initial imbalance of limit orders at the highest bid/lowest ask levels. At our level of statistical errors it appears that the average price change scales approximately linearly with the excess supply (or demand). This plot once again confirms that the influence of the state of the order book on the price is a real and sizable effect. This effect is more pronounced for relatively low-volume stocks such as JDSU and BRCM and less so for a high-volume stock like CSCO. However, as shown in the inset of Fig. 8, normalization of the x -axis by the average number of shares of the stock traded between two of our snapshots, which is proportional to the daily volume, and normalization of the y -axis (1-min price change) by the average daily price (or alternatively the use of returns $r(t) = p(t + \Delta t)/p(t) - 1$), makes all three curves approximately collapse on top of each other.

In conclusion, we have presented an empirical study of statistical properties of a limit order book using the high frequency data collected in the NASDAQ Level II system. It was observed that the distribution of market (or marketable limit) orders has power law tails characterized by an exponent $1 + \mu_{\text{market}} = 2.4 \pm 0.1$. The distribution of limit order sizes is also consistent with a power law with an exponent close to 2. However, it was found that a log-normal distribution provides a better fit to the cumulative distribution of limit order sizes over a wider range. The depth of the order book measured as a virtual price impact of a hypothetical large market order was found to be a non-linear function of its size. This non-linearity is due to the decay in the

² To be precise two prices used to calculate the price change were separated by 20 of our screen snapshots, which are approximately 3 s apart.

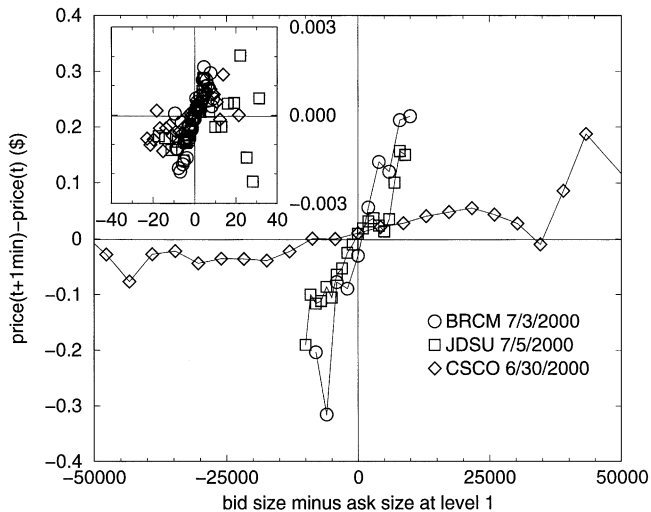


Fig. 8. The average price change one minute after an excess supply or demand was observed as a function of the excess demand, i.e., volume at the highest bid minus volume at the lowest ask. The inset shows the same data but with the y -axis normalized by the average daily price of the stock, and the x -axis normalized by the average volume (number of shares traded) between two snapshots.

density of limit orders (quotes) away from the most recent transaction price. In reality though, this virtual impact is probably much softened by actions of speculators, so that the convexity of the non-linear part may even change its sign. A large imbalance in the number of limit orders at the highest bid and lowest ask sides of the book leads to a predictable average price change which is in accord with intuitive notions regarding supply and demand. This effect seems to disappear at a time scale of several minutes. The short-term average price change linearly depends on the imbalance in the total volume of limit orders at the inside bid and ask prices. These empirical findings may prove to be useful in narrowing down the list of models, used to explain the set of stylized facts about market price fluctuations. Even more importantly, this work may shift the attention of the econophysics community towards more realistic order book based price formation mechanisms. Work is currently underway to add some of the observed empirical features to the simple toy model of order-driven markets proposed by one of us in Ref. [13]. In particular, we plan to check the effect that broad (power law) distributions of limit and market order sizes would have on the critical exponents of this model.

Work at Brookhaven National Laboratory was carried out under Contract No. DE-AC02-98CH10886, Division of Material Science, U.S. Department of Energy.

References

- [1] P. Gopikrishnan, M. Meyer, L.A.N. Amaral, H.E. Stanley, Inverse cubic law for the probability distribution of stock price variations, *Eur. Phys. J. B* 3 (1998) 139.

- [2] M.M. Dacorogna, U.A. Muller, R.J. Nagler, R.B. Olsen, O.V. Pictet, *J. Int. Money Finance* 12 (1993) 413.
- [3] D.M. Guillaume, M.M. Dacorogna, R.D. Dave, U.A. Muller, R.B. Olsen, O.V. Pictet, *Finance and Stochastics* 1 (1997) 95.
- [4] M.M. Dacorogna, U.A. Muller, O.V. Pictet, C.G. de Vries, The distribution of extremal foreign exchange rate returns in extremely large data sets, The Olsen Group Internal Document, UAM 1992-10-22 (1995) (http://www.olsen.ch/library/research/oa_working.html).
- [5] R. Cont, M. Potters, J.-P. Bouchaud, Scaling in stock market data: stable laws and beyond, Lecture given at Les Houches Workshop on Scale Invariance, March 1997 (<http://xxx.lanl.gov/abs/cond-mat/9705087>).
- [6] Y. Liu, P. Gopikrishnan, P. Cizeau, M. Meyer, C.-K. Peng, H.E. Stanley, The statistical properties of the volatility of price fluctuations, *Phys. Rev. E* 60 (1999) 1390.
- [7] Y.-C. Zhang, Modeling market mechanism with evolutionary games, *Europhys. News* 29 (1998) 51.
- [8] D. Challet, M. Marsili, Y.-C. Zhang, Modeling market mechanism with minority game, 1999, cond-mat/9909265.
- [9] D. Challet, M. Marsili, Y.-C. Zhang, Stylized facts of financial markets and market crashes in minority games, 2001, cond-mat/0101326.
- [10] R. Cont, J.-P. Bouchaud, Herd behavior and aggregate fluctuations in financial markets, 1997, cond-mat/9712318.
- [11] R. Cont, J.-P. Bouchaud, *Macroeconomic Dynamics* 4 (2) (2000) 170.
- [12] D. Sornette, D. Stauffer, H. Takayasu, Market fluctuations, multiplicative and percolation models, size effects and predictions, 1999, cond-mat/9909439.
- [13] S. Maslov, Simple model of a limit order-driven market, *Physica A* 278 (2000) 571.
- [14] B. Biais, P. Hilton, C. Spatt, An empirical analysis of the limit order book and the order flow in the Paris Bourse, *J. Finance* 50 (1995) 1655.
- [15] K. Hedvall, J. Niemeyer, G. Rosenqvist, Do buyers and sellers behave similarly in a limit order book? A high-frequency data examination of the Finnish stock exchange, *J. Empirical Finance* 4 (1997) 279, and references therein.
- [16] P. Gopikrishnan, V. Plerou, X. Gabaix, H.E. Stanley, Statistical properties of share volume traded in financial markets, *Phys. Rev. E* 62 (2000) R4493.
- [17] J. Niemeyer, P. Sandas, An empirical analysis of the trading structure at the Stockholm Stock Exchange, Stockholm School of Economics, Working paper No. 44, 1995.
- [18] Y.-C. Zhang, Toward a theory of marginally efficient markets, *Physica A* 269 (1999) 30.